

Resilient Control Systems: Next Generation Design Research

Craig G. Rieger[†], *Senior Member, IEEE*, David I. Gertman[†], Miles. A. McQueen[†], *Member, IEEE*

[†]Idaho National Laboratory, Idaho Falls, Idaho, USA

Abstract — Since digital control systems were introduced to the market more than 30 years ago, the operational efficiency and stability gained through their use have fueled our migration and ultimate dependence on them for the monitoring and control of critical infrastructure. While these systems have been designed for functionality and reliability, a hostile cyber environment and uncertainties in complex networks and human interactions have placed additional parameters on the design expectations for control systems.

Keywords —resilience, control systems, complex networks, cyber security, data fusion.

I. INTRODUCTION

A preeminent objective for corporate and government organizations is state awareness, a comprehensive understanding of security and safety, for critical infrastructures [1]. Unlike the traditional control sense, state in this case confers the information that concerns the ultimate objective of maintaining control. Given the dependence of critical infrastructure on control systems for automation, the integrity of these systems and their ability to provide owner/operators a high degree of state awareness is essential in attaining a high degree of public acceptability. Operators as well as government are therefore burdened to ensure they have a timely understanding of the status of their plant or all plants, respectively, to ensure efficient operations and public protection. This characterization is a significant objective that must consider many aspects of instrumentation, control, and intelligent systems in order to achieve the required result. These aspects include sensory, communication, analysis, decision, and human system interfaces necessary to achieve fusion of data and presentation of results that will provide an understanding of what issues are important and why.

Coupled with the need for state awareness is resilient design, which necessitates a paradigm shift with respect to the methods historically used in control system design. Traditional trust relationships in peer communications are no longer satisfactory since they ignore the malicious actor or actions. While fundamental monitoring and control principles can be applied to achieve a level of success in

preventing security events, these techniques are primarily reactive.

The basis of resilient design requires consideration of all threats and measures by which we determine proper operation. These measures, which can be categorized as cyber and physical security, process efficiency and stability, and process compliancy, provide the operating requirements that are monitored for state awareness and definition of the state space. Traditional concepts of redundancy, diversity, and defense in depth that were once only considered for reliability can be broadened for application to all measures. New concepts that research the human system responses, both benevolent operator and malicious actor interactions, as well as the complex interdependencies of distributed control systems require consideration. The move from reactive to proactive control of plants and mechanisms by which the evaluation and verification of designs is considered all the way from design through implementation stages of resilient control systems (Fig. 1) is enabled by this paradigm shift.




Fig. 1. Resilient Control System

II. DEFINING RESILIENCE

Current research in resilience has focused in two notable areas, organizational and information technology. Organizational resilience considers the ability of an organization to survive in the face of threats, including the prevention or mitigation of unsafe, hazardous or detrimental conditions that threaten its very existence [2]. Information technology resilience considers stability and quality of service in the face of threats to the computing and networking infrastructure. Some might consider control systems that utilize typical information technology components, such as off-the-shelf-computers and IP-based networks, as just a subset of the same. However, control

systems provide a whole layer of complexity not adequately encompassed within the objectives of information technology [3]. The focus of this difference is the key significance that these control systems has in ensuring proper operation of critical infrastructures, including energy and manufacturing facilities.

In considering a definition of resilience, it has been suggested that “Resilient control systems are those that tolerate fluctuations via their structure, design parameters, control structure and control parameters [4].” While this definition is broad, it does not directly consider the presence and necessity of malicious actors. Therefore, another definition might be “an effective reconstitution of control under attack from intelligent adversaries,” which was recently proposed [5]. However, this definition appears to focus only on resiliency in response to the intelligent adversary. True resiliency, however, must consider what represents the proper operation of the process application in the face of many upset conditions, including those attributable to threats from undesirable human interactions. Let’s consider some precepts in the areas of state awareness and resilient design:

- State Awareness
 - Has to be a given for any measure or threat consideration affecting normalcy
 - Must be viable for unexpected threats, and therefore, also those expected
 - Allows supervisory subordinates defined autonomy for a faster control response.
- Resilient Design
 - Comes at a price, and equates to accepted risk given an understanding of consequence
 - Maintains an accepted level of normalcy in the operation of the control system, and as a result, also in the process application
 - In the presence of threat supports mitigation as well as restoration of function.

The word “recovery” was not used in these precepts because its function is assumed to be an underlying premise of resilience. The reasoning for this will be illustrated through a few examples. If resilience is defined in terms of force on a rubber ball, there will be a recovery of the original dimensions once the force is removed. However, if the force exceeds the yield strength of the ball, the ball will be deformed. If resilience is defined in terms of a chemical surge tank, which is placed between coupled processes to prevent instability or shutdowns due to variations in flow, tank levels will rise and fall for variations in flow. Recovery comes in the form of maintaining the desired discharge flow, irrespective of the input flow. However, if the level drops enough to empty the tank or increases above the tank capacity, large variations in the discharge flow will result. With both of these examples a clear limit of resilience is indicated, and normalcy is only achieved if the control system can recover and maintain the system within this limit.

It is worth noting that if the desired level of control system normalcy is achieved without reaching the resilience limit, then the need for state awareness lies only in confirmation of normalcy. However, if conditions change in such a fashion that the bases for resilience limit changes, then recovery cannot be assumed and resilience may not be maintained. It is under these circumstances, therefore, that an awareness of the state is required to enable a response and circumvent the affects of reaching the limit. Phrased another way, state awareness ensures that in the face of any threat, knowledge of unacceptable control system behavior is maintained. With this in mind, the following is therefore proposed as a definition of a resilient control system:

A resilient control system is one that maintains state awareness and an accepted level of operational normalcy in response to disturbances, including threats of an unexpected and malicious nature.

III. AREAS OF RESILIENCE

Several areas are topical to control system resilience. These complement the fundamental concept of dependable or reliable computing by characterizing resilience in regard to the particular control system concerns, including design considerations that provide a level of understanding and assurance in the safe and secure operation of a plant or facility. These areas are presented below with discussion to characterize the basis for consideration as an area of resilience.

A. Human Systems

The human ability to quickly understand novel situations, employ heuristics and analogy can provide additional control system resilience. On the other hand there are situations in which we may have a general inability to reproducibly predict human behavior. This may be true in situations of fatigue or high stress or decision making under high levels of uncertainty. Bayesian methods provide one method by which to take into account evidence regarding human response, but this is one among many approaches. The literature in human reliability analysis provides an orientation regarding ergonomics, workload, complexity, training, experience, etc., which may be used to characterize and quantify human actions and decisions.

Digital technology, used to benefit control system interaction, can from the operators perspective, provide additional complexity. For example, more information can be presented to the human operator to base a response. However, the response could be completely automated, human manipulated, or a combination of both. The dependencies and rules for these complex interactions, or mixed initiative, are not necessarily well defined or clear. Resiliency results from understanding of this complexity, ensuring through human factor and design an error tolerant control system results that complements perception, fusion, and decision making.

B. Complex Networks

As control systems become more decentralized, the ability to characterize interactions, performance and security becomes more critical to ensuring resilience. While more decentralization can provide additional reliability due to implicit redundancy and diversity, it may also provide more avenues or vectors to cyber attack. Therefore, the design of complex networks needs to consider all factors that influence resilience, and optimize for multiple considerations [6].

Global stability is often perceived as something that can be achieved by local minimization of all process unit operations, many of which are contained in a facility. However, there is no assurance that global stability can be achieved in this manner, and in addition, this philosophy maintains a reactionary control paradigm by its nature. However, considering the latencies in digital control systems, there is a tendency as well as a desire to provide faster responses when the feedback and response occur close to the point of interaction with the application. Therefore, it is suggested that a true global optimization coupled with a local interaction can achieve both the assurance of a global minima, and an acceptable response when designing control system architecture.

C. Cyber Awareness

Because of the human element of a malicious actor, traditional methods of achieving reliability cannot be used to characterize cyber awareness and resilience. Dynamic mechanisms of probabilistic risk analysis that can link human reliability with the system state are still maturing. The intellectual level and background of the adversary makes stochastic methods unusable due to the variability of both the objective and the motives. In addition, the strength of the adversary is increased because the existing control system architecture is not random, and response characteristics are reproducible. Therefore, a resilient design can find strength in similar fashion by becoming atypical of normal control system architectural design, and appearing random in response and characteristics to the adversary.

Characterization of health or wellness from a cyber perspective is purely empirical, as prediction of the future is based on past events. While there are barriers in place to exclude known types of adversarial communication, state awareness cannot be assured because of the limited availability of diverse sensing. Determination of the actual cause of an abnormal event can only occur only after forensics are completed. Patterns or routines are analyzed and are used to provide comparisons to understand anomalies. However, while this understanding provides an interesting perspective, it may be very limited in predicting future behavior of the adversary.

D. Data Fusion

The nature of the various data types associated with proper operation or performance of critical infrastructure, including cyber and physical security, process efficiency and stability, and process compliancy is diverse. How these data are consumed to generate information will help

determine whether appropriate judgments are made, whether by automated and/or human mechanisms. There are several issues that are addressed by data fusion, including the following ones:

- Reduction - The reduction of data to provide only that information necessary for the human or automation scheme to provide the appropriate response, i.e., to prevent a common issue of information overload.
- Identification - Validation and invalidation of causes for events, e.g., a process upset is due to a failed valve and not a cyber attack.
- Improved characterization and knowledge - Development of new information that helps to better characterize the process application, e.g., mining of process temperatures along with process flows provides a better interpretation of stability.

While many of the techniques required to perform data fusion are well known, application to the diverse types of data represented within the measures of performance provide a distinct challenge [7]. This is nowhere more evident than the fusion of cyber and process data to not only indicate whether an event is cyber specific, whether due to an adversary or network problem, or actually represents a process upset. The effort to address this situation could be split into two parts: i) developing the appropriate data to characterize the cyber threat, and ii) combining the spatial and temporal aspects of both process and cyber data to confirm the cause of the process upset.

IV. RESEARCH AREAS

In considering the areas of resilience, two overall categories, state awareness and resilient design, which were given during the development of the definition, can be used to cover both the measures of performance and the actual design. A brief explanation of each is given below.

A. State awareness

In defining state awareness, one must reflect on the fundamental reasons for installing a control system in the first place. From a monitoring standpoint, these control systems are expected to provide a sufficient knowledge of operating parameters that represent a basis for decisions. However, there are a number of measures that are based on the uses of the data, which also provide the basis for establishing performance requirements. From the smallest to the largest control system, maintaining a state awareness of everything that can affect its normalcy must be performed. These measures have previously been identified as cyber and physical security, process efficiency and stability, and process compliancy.

However, gaining state awareness is more than having all the appropriate sources of data. What the consumer of the data really requires is the information necessary to maintain the normalcy of the control system, within the limits of authority that have been provided. This requires focusing and prioritizing information based upon an intelligent fusion of data. Intelligent fusion not only

reduces the level of information provided to the consumer, but also generates data better characterizing the awareness state space via observers and predictors.

B. Resilient design

Resilient control system design complements traditional considerations of reliability and dependable computing, which are well established research areas. However, while reliability design brings with it the fundamental considerations of platform operations and communications with no particular focus on the use of the platform, resilient design must consider the attributes that are particular to control systems. Resilient design provides a paradigm shift on how we look at control system design, where traditional redundancy would have been implemented based on a particular vendor's perspective on reliability design. These designs find basis in the characterization of reproducible and understood events, and while applied to control systems, similar concepts could equally be applied to many types of microprocessor-based applications.

The concepts of safety instrumented systems have taken a step toward considering elements of control system design that are unique to the process application. For example, the control system and its function to prevent unsafe conditions in the process application are considered when determining probability of failure. In a traditional sense, component failure alone was the concern. However, to be resilient, there are notional ideas that come from the areas of resilience already discussed. In human systems, the unpredictability of the human threatens resilience, while the innate ability to adapt reinforces resilience. Similarly with cyber awareness, the unpredictability of the attacker threatens resilience; however, in this case the ability to adapt is also an added threat. With complex networks, latencies and disruptions in communications may affect the stability of coupled control loops, negatively impacting the resilience. These threats to resilience, considered specifically in regard to desired operation of the process application, form the paradigm under which resilience in the context of this paper and research finds its basis.

In providing monitoring and control capabilities, the basic element of a control system is its underlying feedback control loop, which may be hosted on many communicating platforms, including transmitters, converters, logic solvers, and operator displays. How these elements are build into an integrating architecture can vary, especially when considering next generation resilient designs. In identifying the best method, however, the considerations of complex network design are necessary to build and optimize the interactions of the various elements. When the human elements are considered within this architecture, the purpose of data fusion can be realized. Data fusion is normally considered a method to concentrate or combine data to yield information and knowledge, which in this case provides state awareness and the basis for decisions. However, while the principles of data fusion as described provide a more focused perspective to provide more resilience to the friendly human, by the nature these principles can also be

“reversed,” so to speak, to provide the contrary results. It is this perspective that is needed to counteract the negative impact on resilience brought by the malicious actor trying to undermine a control system. Therefore, it is desired to increase, not decrease, the confusion of the malicious actor by undermining his understanding of the control system.

V. TASK AREAS

In considering the research areas of state awareness and resilient design for resilient control systems, several task areas become evident as identified within Fig. 2-3. The task areas within the state awareness research area have been previously defined as cyber and physical security, process efficiency and stability, and process compliancy. Within these task areas are formed the requirements that define measures of normalcy in the context of a process application. Within each of the state awareness task areas, data is retrieved, analyzed, fused and tailored to the user or consumer of the data, whether operator, manager or engineer. There are a few reasons why this procedure is necessary. Within process operations the availability of data is often in excess of what is needed to maintain normalcy, but is prevalent and given in quantities that overwhelm the consumer. However with cyber awareness, data can be insufficient to ensure normalcy, even if provided in overwhelming quantities. This follows as process sensing and measurement is an old field and cyber security is a relatively new field; hence, is not as mature in characterization.

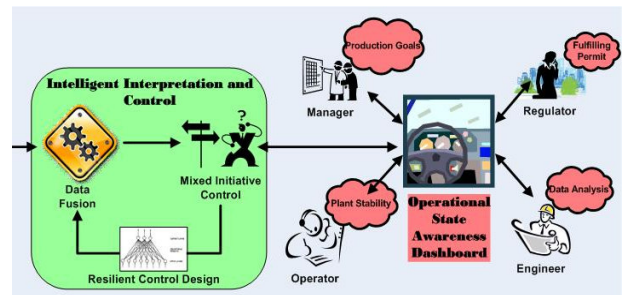


Fig. 2. State Awareness

Within the resilient design area, task areas define the appropriate control responses to maintain normalcy based on the state of the process application. These control responses consider both the human and automation mechanisms, and mixed initiative. Although called automation, which could be characterized as regimented, the preferred methods are often a combination of techniques. Regimented responses can be termed as hard computing, as they are based on “hard” physics and first principles models of the process, however simplified. With soft computing, responses are based on methods that attempt to capture the ability of the human or biological systems to adapt and provide a “soft” response based on the environmental observations. Combinations of these techniques are necessary when considering the resilience of complex processes and distributed networks. In addition, just as state awareness considers all aspects that may affect normalcy, the application of control must

consider not only local requirements but also top down perspectives. Considering a related example, any organization will fail that does not have a governing strategy or organizational leadership. Similarly, control designs cannot be implemented as a loose collection of individual control loops and expect to ensure the stability, efficiency and security of the whole facility. A supervisory design is necessary to provide the required oversight of the facility, providing a hierarchy of both responsibility and autonomy that ensures normalcy is maintained. Each level of the hierarchy can use a combination of soft and hard computing techniques, providing intelligent interpretation and control.

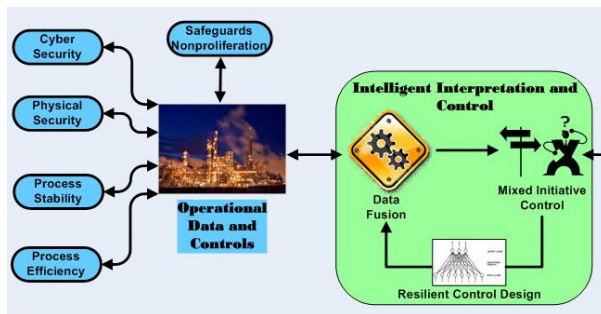


Fig. 3. Resilient Design

VI. CONCLUSION

The discussion presented provides a conceptual framework and brief overview of the architectural considerations of the control system application and provides a basis for resilient control system research. While digital control has afforded the opportunity to remotely sense and control processes for several decades, the complexity of the control systems and the facilities they are intended to safely and reliably control has increased. Cyber connectivity, which provides a fairly inexpensive avenue to infiltrate control systems, has increased the complexity. A better understanding and resulting optimization will require a mathematical representation of this complexity. There is little tolerance in our society for loss or failure of infrastructure, equating to loss of normalcy in the routines of the populace, as a minimum, or loss of life as a maximum. Many events can initiate these failures, natural or man-made, and there will be a price for the level of resilience desired. However, the mechanisms of resilient control systems, state awareness and resilient designs, promise to provide a logical vehicle to respond to these events.

REFERENCES

- [1] Critical Infrastructure Protection, Department of Homeland Security Faces Challenges in Fulfilling Cybersecurity Responsibilities, GAO-05-434, May 2005.
- [2] E. Hollnagel, D. D. Woods, and N. Leveson, *Resilience Engineering: Concepts and Precepts*, Ashgate Publishing, Aldershot Hampshire, UK, 2006.
- [3] F.-Y. Wang and D. Liu, *Networked Control Systems: Theory and Applications*, Springer-Verlag, London, UK, 2008.

- [4] S. M. Mitchell and M. S. Mannan, "Designing Resilient Engineered Systems," *Chemical Engineering Progress*, Vol. 102, No. 4, pp. 39-45, April 2006.
- [5] Proceedings of the 1st International Symposium on Resilient Control Systems, Idaho Falls, ID, 2008.
- [6] S.P. Meyn, *Control Techniques for Complex Networks*, Cambridge University Press, New York, NY, 2008.
- [7] H. B. Mitchell, *Multi-Sensor Data Fusion*, Springer-Verlag, Berlin, 2007.